

PORTAL

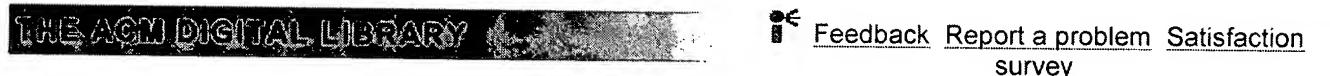
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+remote +direct +memory +access +iscsi

SEARCH

Terms used: **remote direct memory access iscsi**

Found 29 of 209,709

Sort results by

relevance Save results to a Binder

Try an Advanced Search

Display results

expanded form Search Tips

Try this search in The ACM Guide

 Open results in a new window

Results 1 - 20 of 29

Result page: **1** [2](#) [next](#)

Relevance scale

1 Storage protocol designs: A study of iSCSI extensions for RDMA (iSER)

Mallikarjun Chadalapaka, Hemal Shah, Uri Elzur, Patricia Thaler, Michael Ko
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications NICELI '03**

Publisher: ACM Press

Full text available: pdf(281.32 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The iSCSI protocol is the IETF standard that maps the SCSI family of application protocols onto TCP/IP enabling convergence of storage traffic on to standard TCP/IP fabrics. The ability to efficiently transfer and place the data on TCP/IP networks is crucial for this convergence of the storage traffic. The iWARP protocol suite provides Remote Direct Memory Access (RDMA) semantics over TCP/IP networks and enables efficient memory-to-memory data transfers over an IP fabric. This paper studies the ...

Keywords: DA, DDP, DI, Datamover, MPA, RDMA, RDMAP, SCSI, Verbs, iSCSI, iSER, iWARP

2 Promises and reality: Server I/O networks past, present, and future

Renato John Recio
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications NICELI '03**

Publisher: ACM Press

Full text available: pdf(225.62 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Enterprise and technical customers place a diverse set of requirements on server I/O networks. In the past, no single network type has been able to satisfy all of these requirements. As a result several fabric types evolved and several interconnects emerged to satisfy a subset of the requirements. Recently several technologies have emerged that enable a single interconnect to be used as more than one fabric type. This paper will describe the requirements customers place on server I/O networks; t ...

Keywords: 10 GigE, Cluster, Cluster Networks, Gigabit Ethernet, I/O Expansion Network, IOEN, InfiniBand, LAN, PCI, PCI Express, RDMA, RNIC, SAN, Socket Extensions, TOE, iONIC, iSCSI, iSER

3 Scalability: Direct-pNFS: scalable, transparent, and versatile access to parallel file

 **systems**

Dean Hildebrand, Peter Honeyman

June 2007 **Proceedings of the 16th international symposium on High performance distributed computing HPDC '07**

Publisher: ACM Press

Full text available:  pdf(397.10 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Grid computations require global access to massive data stores. To meet this need, the GridNFS project aims to provide scalable, high-performance, transparent, and secure wide-area data management as well as a scalable and agile name space.

While parallel file systems give high I/O throughput, they are highly specialized, have limited operating system and hardware platform support, and often lack strong security mechanisms. Remote data access tools such as NFS and GridFTP overcome some ...

Keywords: distributed file system, nFSv4, pNFS, parallel I/O

4 Direct Cache Access for High Bandwidth Network I/O

 Ram Huggahalli, Ravi Iyer, Scott Tetruck

May 2005 **ACM SIGARCH Computer Architecture News , Proceedings of the 32nd annual international symposium on Computer Architecture ISCA '05**, Volume 33 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  pdf(194.52 KB) Additional Information: [full citation](#), [abstract](#), [cited by](#), [index terms](#)

Recent I/O technologies such as PCI-Express and 10Gb Ethernet enable unprecedented levels of I/O bandwidths in mainstream platforms. However, in traditional architectures, memory latency alone can limit processors from matching 10 Gb inbound network I/O traffic. We propose a platform-wide method called Direct Cache Access (DCA) to deliver inbound I/O data directly into processor caches. We demonstrate that DCA provides a significant reduction in memory latency and memory bandwidth for receive in ...

5 Storage protocol designs: NFS over RDMA

 Brent Callaghan, Theresa Lingutla-Raj, Alex Chiu, Peter Staubach, Omer Asad

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications NICELI '03**

Publisher: ACM Press

Full text available:  pdf(126.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

The NFS filesystem was designed as a work-group filesystem, making a central file store available to and shared between a number of client workstations. However, more recently NFS has grown in popularity in the server room, connecting large application servers with back-end file servers. In this environment, where high-speed access to data is critical, high capacity interconnects like gigabit Ethernet, Fibre Channel and Infiniband are to be expected. With RDMA technology we can fully utilize the ...

6 I/O: miNI: reducing network interface memory requirements with dynamic handle lookup

 Reza Azimi, Angelos Bilas

June 2003 **Proceedings of the 17th annual international conference on Supercomputing ICS '03**

Publisher: ACM Press

Full text available:  pdf(289.75 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recent work in low-latency, high-bandwidth communication systems has resulted in building user-level Network Interface Controllers (NICs) and communication abstractions

that support direct access from the NIC to applications virtual memory to avoid both data copies and operating system intervention. Such mechanisms require the ability to directly manipulate user--level communication buffers for delivering data and achieving protection. To provide such abilities, NICs must maintain appropriate t ...

Keywords: parallel architectures, system area networks

7 Integrated network interfaces for high-bandwidth TCP/IP

 Nathan L. Binkert, Ali G. Saidi, Steven K. Reinhardt
October 2006 **ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , Proceedings of the 12th international conference on Architectural support for programming languages and operating systems ASPLOS-XII**, Volume 40 , 41 , 34 Issue 5 , 11 , 5

Publisher: ACM Press

Full text available:  pdf(420.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper proposes new network interface controller (NIC) designs that take advantage of integration with the host CPU to provide increased flexibility for operating system kernel-based performance optimization. We believe that this approach is more likely to meet the needs of current and future high-bandwidth TCP/IP networking on end hosts than the current trend of putting more complexity in the NIC, while avoiding the need to modify applications and protocols. This paper presents two such NICs ...

Keywords: TCP/IP performance, interfaces, network, zero-copy

8 Features: TCP Offload to the Rescue

 Andy Currid
May 2004 **Queue**, Volume 2 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.82 MB)
 html(28.93 KB) Additional Information: [full citation](#), [index terms](#)

9 Promises and reality: Performance measurements of a user-space DAFS server with a database workload

 Samuel A. Fineberg, Don Wilson
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications NICELI '03**

Publisher: ACM Press

Full text available:  pdf(366.48 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We evaluate the performance of a user-space Direct Access File System (DAFS) server and Oracle Disk Manager (ODM) client using two synthetic test codes as well as the Oracle database. Tests were run on 4-processor Intel Xeon-based systems running Windows 2000. The systems were connected with ServerNet II, a Virtual Interface Architecture (VIA) compliant system area network. We compare the performance of DAFS/ODM and local-disk based I/O, measuring I/O bandwidth and latency. We also compare the r ...

Keywords: DAFS, Database, File Systems, I/O, Networks, Performance Evaluation, RDMA

10 Novel approaches: Engineering a user-level TCP for the CLAN network

 Kieran Mansley

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications NICELI '03**

Publisher: ACM Press

Full text available:  pdf(142.75 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

As networks and I/O systems converge and the bandwidth of networks increases, conventional approaches to networking are struggling to deliver the performance and flexibility required. CLAN (Collapsed LAN) is a high performance user-level network targeted at the server room. It supports RDMA and programmed I/O (PIO). We have implemented a set of IP based protocols at user level, and shown how true zero copy transmission (without modifying the sockets API) and reception can be achieved. In this paper ...

11 A study of the impact of direct access I/O on relational database management systems

Heidi Scott, Patrick Martin, Berni Schiefer

September 2002 **Proceedings of the 2002 conference of the Centre for Advanced Studies on Collaborative research CASCON '02**

Publisher: IBM Press

Full text available:  pdf(117.37 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Direct access I/O allows an application program to send requests directly to the I/O subsystem without involving the operating system. We believe that data-intensive applications such as database management systems (DBMSs) stand to reap significant performance benefits from direct access I/O. In this paper we describe an initial attempt to verify this claim. We present a prototype direct access file system and describe a set of experiments we conducted with the prototype and a modified ve ...

12 Scalability: High performance and scalable I/O virtualization via self-virtualized devices

 Himanshu Raj, Karsten Schwan

June 2007 **Proceedings of the 16th international symposium on High performance distributed computing HPDC '07**

Publisher: ACM Press

Full text available:  pdf(420.44 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

While industry is making rapid advances in system virtualization, for server consolidation and for improving system maintenance and management, it has not yet become clear how virtualization can contribute to the performance of high end systems. In this context, this paper addresses a key issue in system virtualization - how to efficiently virtualize I/O subsystems and peripheral devices. We have developed a novel approach to I/O virtualization, termed *self-virtualized devices*, which i ...

Keywords: virtual devices, virtualization

13 Experiences with VI communication for database storage

 Yuanyuan Zhou, Angelos Bilas, Suresh Jagannathan, Cezary Dubnicki, James F. Philbin, Kai Li May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  pdf(1.29 MB)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper examines how VI-based interconnects can be used to improve I/O path performance between a database server and the storage subsystem. We design and implement a software layer, DSA, that is layered between the application and VI. DSA takes advantage of specific VI features and deals with many of its shortcomings. We provide and evaluate one kernel-level and two user-level implementations of DSA. These implementations trade transparency and generality for performance at different degrees ...

Keywords: Storage system, cluster-based storage, Database storage, storage area network, User-level Communication, Virtual Interface Architecture, processor overhead

14 Data Reservoir: utilization of multi-gigabit backbone network for data-intensive research

Kei Hiraki, Mary Inaba, Junji Tamatsukuri, Ryutaro Kurusu, Yukichi Ikuta, Hisashi Koga, Akira Zinzaki

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing Supercomputing '02**

Publisher: IEEE Computer Society Press

Full text available:  pdf(289.69 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We propose data sharing facility for data intensive scientific research, "Data Reservoir"; which is optimized to transfer huge amount of data files between distant places fully utilizing multi-gigabit backbone network. In addition, "Data Reservoir" can be used as an ordinary UNIX server in local network without any modification of server softwares. We use low-level protocol and hierarchical striping to realize (1) separation of bulk data transfer and local accesses by caching, (2) file-system tr ...

15 Runtime systems: Optimization and bottleneck analysis of network block I/O in commodity storage systems

Manolis Marazakis, Vassilis Papaefstathiou, Angelos Bilas

June 2007 **Proceedings of the 21st annual international conference on Supercomputing ICS '07**

Publisher: ACM Press

Full text available:  pdf(442.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Building commodity networked storage systems is an important architectural trend; Commodity servers hosting a moderate number of consumer-grade disks and interconnected with a high-performance network are an attractive option for improving storage system scalability and cost-efficiency. However, such systems incur significant overheads and are not able to deliver to applications the available throughput. We examine in detail the sources of overheads in such systems, using a working prototype ...

Keywords: I/O performance optimization, RDMA, block-level I/O, commodity servers

16 I/O--communication: Large files, small writes, and pNFS

Dean Hildebrand, Lee Ward, Peter Honeyman

June 2006 **Proceedings of the 20th annual international conference on Supercomputing ICS '06**

Publisher: ACM Press

Full text available:  pdf(2.28 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Workload characterization studies highlight the prevalence of small and sequential data requests in scientific applications. Parallel file systems excel at large data transfers but sometimes at the expense of small I/O performance. pNFS is an NFSv4.1 high-performance enhancement that provides direct storage access to parallel file systems

while preserving NFSv4 operating system and hardware platform independence. This paper demonstrates that distributed file systems can increase write throughput ...

Keywords: NFSv4, distributed file system, pNFS, parallel I/O, parallel file system, small write performance improvement

17 Network attached storage architecture

 Garth A. Gibson, Rodney Van Meter
November 2000 **Communications of the ACM**, Volume 43 Issue 11

Publisher: ACM Press

Full text available:  pdf(224.67 KB)

 html(43.39 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 Short papers -- works in progress: Toward a threat model for storage systems

 Ragib Hasan, Suvda Myagmar, Adam J. Lee, William Yurcik
November 2005 **Proceedings of the 2005 ACM workshop on Storage security and survivability StorageSS '05**

Publisher: ACM Press

Full text available:  pdf(258.24 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The growing number of storage security breaches as well as the need to adhere to government regulations is driving the need for greater storage protection. However, there is the lack of a comprehensive process to designing storage protection solutions.

Designing protection for storage systems is best done by utilizing proactive system engineering rather than reacting with ad hoc countermeasures to the latest attack du jour. The purpose of threat modeling is to organize system threats and vulnera ...

Keywords: security, storage system, threat model

19 Features: Storage Systems: Not Just a Bunch of Disks Anymore

 June 2003 **Queue**, Volume 1 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.29 MB)

 html(31.84 KB)

Additional Information: [full citation](#), [index terms](#)

20 Separating Abstractions from Resources in a Tactical Storage System

Douglas Thain, Sander Klous, Justin Wozniak, Paul Brenner, Aaron Striegel, Jesus Izaguirre
November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

Publisher: IEEE Computer Society

Full text available:  pdf(401.40 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Sharing data and storage space in a distributed system remains a difficult task for ordinary users, who are constrained to the fixed abstractions and resources provided by administrators. To remedy this situation, we introduce the concept of a tactical storage system (TSS) that separates storage abstractions from storage resources, leaving users free to create, reconfigure, and destroy abstractions as their needs change. In this paper, we describe how a TSS can provide a variety of filesystem an ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) | [Purchase History](#) | [Cart](#) |

Welcome United States Patent and Trademark Office

Search Results**BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "((remote<in>metadata) <and> (direct<in>metadata))<and> (memory<in>meta... "

Your search matched 68 of 1634821 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.[e-mail](#)**» Search Options**[View Session History](#)[New Search](#)**» Key**

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard

Modify Search ((remote<in>metadata) <and> (direct<in>metadata))<and> (memory<in>meta...) Check to search only within this results setDisplay Format: Citation Citation & Abstract

View: 1-

1. Memory Management Support for Multi-Programmed Remote Direct Mem (RDMA) Systems

Magoutis, K.;

[Cluster Computing, 2005. IEEE International](#)

Sept. 2005 Page(s):1 - 8

Digital Object Identifier 10.1109/CLUSTR.2005.347031

[AbstractPlus](#) | Full Text: [PDF\(371 KB\)](#) IEEE CNF[Rights and Permissions](#)**2. FRAMP: a fast remote direct memory access and message passing netw**

Gang Shi; Mingchang Hu; Hongda Yin; Weiwu Hu; Zhimin Tang;

[Communications and Information Technology, 2004. ISCIT 2004. IEEE Internation](#)

on Volume 1, 26-29 Oct. 2004 Page(s):591 - 596 vol.1

Digital Object Identifier 10.1109/ISCIT.2004.1412914

[AbstractPlus](#) | Full Text: [PDF\(755 KB\)](#) IEEE CNF[Rights and Permissions](#)**3. A shared virtual memory network with fast remote direct memory access passing**

Gang Shi; Mingchang Hu; Hongda Yin; Weiwu Hu; Zhimin Tang;

[Cluster Computing, 2004 IEEE International Conference on](#)

20-23 Sept. 2004 Page(s):495

Digital Object Identifier 10.1109/CLUSTR.2004.1392660

[AbstractPlus](#) | Full Text: [PDF\(409 KB\)](#) IEEE CNF[Rights and Permissions](#)**4. Message-based efficient remote memory access on a highly parallel com**

Kodama, Y.; Sakane, H.; Sato, M.; Sakai, S.; Yamaguchi, Y.;

[Parallel Architectures, Algorithms and Networks, 1994. \(ISPAN\) International S](#)

14-16 Dec. 1994 Page(s):135 - 142

Digital Object Identifier 10.1109/ISPAN.1994.367154

[AbstractPlus](#) | Full Text: [PDF\(448 KB\)](#) IEEE CNF[Rights and Permissions](#)**5. Evaluating InfiniBand performance with PCI Express**

Jiuxing Liu; Mamidala, A.; Vishnu, V.; Panda, D.K.;

[Micro. IEEE](#)

Volume 25, Issue 1, Jan.-Feb. 2005 Page(s):20 - 29
Digital Object Identifier 10.1109/MM.2005.9

[AbstractPlus](#) | Full Text: [PDF\(160 KB\)](#) IEEE JNL
[Rights and Permissions](#)

-  **6. Out of order incremental CRC computation**
Satran, J.; Sheinwald, D.; Shimony, I.;
[Computers, IEEE Transactions on](#)
Volume 54, Issue 9, Sept. 2005 Page(s):1178 - 1181
Digital Object Identifier 10.1109/TC.2005.151
[AbstractPlus](#) | Full Text: [PDF\(216 KB\)](#) IEEE JNL
[Rights and Permissions](#)
-  **7. Techniques for compiler-directed cache coherence**
Choi, L.; Hock-Beng Lim; Pen-Chung Yew;
[Parallel & Distributed Technology: Systems & Applications; IEEE \[see also IEE](#)
Volume 4, Issue 4, Winter 1996 Page(s):23 - 34
Digital Object Identifier 10.1109/88.544434
[AbstractPlus](#) | References | Full Text: [PDF\(2364 KB\)](#) IEEE JNL
[Rights and Permissions](#)
-  **8. 10-Gigabit iWARP Ethernet: Comparative Performance Analysis with Infir Myrinet-10G**
Rashti, M.J.; Afsahi, A.;
[Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE Inte](#)
26-30 March 2007 Page(s):1 - 8
Digital Object Identifier 10.1109/IPDPS.2007.370480
[AbstractPlus](#) | Full Text: [PDF\(325 KB\)](#) IEEE CNF
[Rights and Permissions](#)
-  **9. RNIC-PI: The last step in standardizing RDMA**
VelurEunni, R.;
[Cluster Computing, 2005. IEEE International](#)
Sept. 2005 Page(s):1 - 7
Digital Object Identifier 10.1109/CLUSTR.2005.347030
[AbstractPlus](#) | Full Text: [PDF\(7657 KB\)](#) IEEE CNF
[Rights and Permissions](#)
-  **10. High Performance Block I/O for Global File System (GFS) with InfiniBand**
Shuang Liang; Weikuan Yu; Panda, D.K.;
[Parallel Processing, 2006. ICPP 2006. International Conference on](#)
Aug. 2006 Page(s):391 - 398
Digital Object Identifier 10.1109/ICPP.2006.47
[AbstractPlus](#) | Full Text: [PDF\(322 KB\)](#) IEEE CNF
[Rights and Permissions](#)
-  **11. Experiences from Debugging a PCIX-based RDMA-capable NIC**
Marazakis, M.; Papaefstathiou, V.; Kalokairinos, G.; Bilas, A.;
[Cluster Computing, 2006 IEEE International Conference on](#)
25-28 Sept. 2006 Page(s):1 - 10
Digital Object Identifier 10.1109/CLUSTR.2006.311919
[AbstractPlus](#) | Full Text: [PDF\(12605 KB\)](#) IEEE CNF
[Rights and Permissions](#)
-  **12. Exploiting RDMA operations for Providing Efficient Fine-Grained Resource Management in Cluster-based Servers**
Vaidyanathan, K.; Jin, H.-W.; Panda, D.K.;
[Cluster Computing, 2006 IEEE International Conference on](#)
25-28 Sept. 2006 Page(s):1 - 10

Digital Object Identifier 10.1109/CLUSTR.2006.311916

[AbstractPlus](#) | Full Text: [PDF\(1692 KB\)](#) IEEE CNF
[Rights and Permissions](#)

- 13. Initial Performance Evaluation of the NetEffect 10 Gigabit iWARP Adapter**
Dalessandro, D.; Wyckoff, P.; Montry, G.;
[Cluster Computing, 2006 IEEE International Conference on](#)
25-28 Sept. 2006 Page(s):1 - 7
Digital Object Identifier 10.1109/CLUSTR.2006.311915
[AbstractPlus](#) | Full Text: [PDF\(188 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 14. An automated approach to improve communication-computation overlap**
Fishgold, L.; Danalis, A.; Pollock, L.; Swany, M.;
[Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International](#)
25-29 April 2006 Page(s):7 pp.
Digital Object Identifier 10.1109/IPDPS.2006.1639590
[AbstractPlus](#) | Full Text: [PDF\(136 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 15. Breaking the connection: RDMA deconstructed**
Rajeev Sivaram; Govindaraju, R.K.; Hochschild, P.; Blackmore, R.; Piyush Chandra;
[High Performance Interconnects, 2005. Proceedings. 13th Symposium on](#)
17-19 Aug. 2005 Page(s):36 - 42
Digital Object Identifier 10.1109/CONECT.2005.9
[AbstractPlus](#) | Full Text: [PDF\(176 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 16. Distributed queue-based locking using advanced network features**
Devulapalli, A.; Wyckoff, P.;
[Parallel Processing, 2005. ICPP 2005. International Conference on](#)
14-17 June 2005 Page(s):408 - 415
Digital Object Identifier 10.1109/ICPP.2005.34
[AbstractPlus](#) | Full Text: [PDF\(200 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 17. Employing an RDMA-based file system for high performance computing**
Velusamy, V.; Skjellum, A.; Kanevsky, A.;
[Networks, 2004. \(ICON 2004\). Proceedings. 12th IEEE International Conference on](#)
Volume 1, 16-19 Nov. 2004 Page(s):66 - 70 vol.1
Digital Object Identifier 10.1109/ICON.2004.1409089
[AbstractPlus](#) | Full Text: [PDF\(701 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 18. Are processors free? Impact on RDMA and protocol off-load technologies**
Krause, M.;
[Network Computing and Applications, 2004. \(NCA 2004\). Proceedings. Third International Symposium on](#)
2004 Page(s):203
Digital Object Identifier 10.1109/NCA.2004.1347778
[AbstractPlus](#) | Full Text: [PDF\(195 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 19. Studying network protocol offload with emulation: approach and preliminary results**
Westrelin, R.; Fugier, N.; Nordmark, E.; Kunze, K.; Lemoine, E.;
[High Performance Interconnects, 2004. Proceedings. 12th Annual IEEE Symposium on](#)
25-27 Aug. 2004 Page(s):84 - 90
Digital Object Identifier 10.1109/CONECT.2004.1375208

[AbstractPlus](#) | Full Text: [PDF\(803 KB\)](#) IEEE CNF
[Rights and Permissions](#)

- 20. High performance MPI-2 one-sided communication over InfiniBand**
Weihang Jiang; Jiuxing Liu; Hyun-Wook Jin; Panda, D.K.; Gropp, W.; Thakur, [Cluster Computing and the Grid, 2004. CCGrid 2004. IEEE International Symposium on](#)
19-22 April 2004 Page(s):531 - 538
Digital Object Identifier 10.1109/CCGrid.2004.1336648
[AbstractPlus](#) | Full Text: [PDF\(590 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 21. Design and implementation of MPICH2 over InfiniBand with RDMA support**
Liu, J.; Jiang, W.; Wyckoff, P.; Panda, D.K.; Ashton, D.; Buntinas, D.; Gropp, V [Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International](#)
26-30 April 2004 Page(s):16
Digital Object Identifier 10.1109/IPDPS.2004.1302922
[AbstractPlus](#) | Full Text: [PDF\(1402 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 22. RDMA control support for fine-grain parallel computations**
Smyk, A.; Tudruj, M.; [Parallel, Distributed and Network-Based Processing, 2004. Proceedings. 12th Conference on](#)
11-13 Feb. 2004 Page(s):208 - 215
Digital Object Identifier 10.1109/EMPDP.2004.1271447
[AbstractPlus](#) | Full Text: [PDF\(556 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 23. A performance analysis of the sockets direct protocol (SDP) with asynchronous 4X InfiniBand**
Cohen, A.; [Performance, Computing, and Communications, 2004 IEEE International Conference on](#)
2004 Page(s):241 - 246
Digital Object Identifier 10.1109/PCCC.2004.1394991
[AbstractPlus](#) | Full Text: [PDF\(718 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 24. Throughput/channel capacity improvement issues for LEO-satellite based sensing systems**
Matar, M.A.; [Recent Advances in Space Technologies, 2003. RAST '03. International Conference on](#)
[Proceedings of](#)
20-22 Nov. 2003 Page(s):488 - 494
Digital Object Identifier 10.1109/RAST.2003.1303965
[AbstractPlus](#) | Full Text: [PDF\(525 KB\)](#) IEEE CNF
[Rights and Permissions](#)
- 25. Micro-benchmark level performance comparison of high-speed cluster interconnects**
Jiuxing Liu; Chandrasekaran, B.; Weikuan Yu; Jiesheng Wu; Buntinas, D.; Kin Panda, D.K.; [High Performance Interconnects, 2003. Proceedings. 11th Symposium on](#)
20-22 Aug. 2003 Page(s):60 - 65
[AbstractPlus](#) | Full Text: [PDF\(269 KB\)](#) IEEE CNF
[Rights and Permissions](#)

View: 1-



© Copyright 2006 IEEE -

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [Gmail](#) [more ▾](#)

[Sign in](#)

Google

remote direct memory access iscsi

[Search](#)

[Advanced Search](#)

[Preferences](#)

New! View and manage your web history

[Web](#)

Results 1 - 10 of about 246,000 for **remote direct memory access iscsi**. (0.18 seconds)

[CommsDesign - iWARP: Reducing Ethernet Overhead in Data Center Designs](#)

Figure 2: RDMA enables the ECA to support **direct memory** reads from/writes to ... unit (MMU) state associated with the application's **memory access** rights. ...

www.commsdesign.com/design_corner/showArticle.jhtml?articleID=51202855 - 46k - Cached - Similar pages

[Architectural Specifications for RDMA over TCP/IP](#)

Sockets Direct Protocol (SDP), **iSCSI** Extensions for RDMA (iSER) and Datamover Architecture for **iSCSI** (DA) Version 1.0 Release Specifications are Complete ...

www.rdmaconsortium.org/ - 9k - Cached - Similar pages

[Remote Direct Memory Access Resources on TechRepublic](#)

The paper considers optimizations to **Remote Procedure Call (RPC)**-based data transfer using either **Remote Direct Memory Access** ... search.techrepublic.com.com/search/Remote+Direct+Memory+Access.html - 50k -

Cached - Similar pages

[Red Hat Readies RHEL 5 for March 14 Launch - Computer Business Review](#)

... the company's high availability clustering software, as well as support for **iSCSI** disk arrays, InfiniBand with **Remote Direct Memory Access** (RDMA), ... www.cbronline.com/article_news.asp?guid=BAA762B7-84F3-48D0-93B1-18CDC979F9BB - 28k - Cached - Similar pages

[\[PPT\] Remote Direct Memory Access \(RDMA\)](#)

File Format: Microsoft Powerpoint - [View as HTML](#)

1999; Mallickarjun Chadalapka , Michael KO , Uri Elzur "A study of **iSCSI** Extension for RDMA (iSER). " "Understanding **Remote Direct Access Memory**". ... www.tntech.edu/ece/reu/2005summer/Presentations/RDMAiSCSI.ppt - Similar pages

[Workshop on Remote Direct Memory Access](#)

Remote Direct Memory Access (RDMA) enables transfer of data across a network ... Applications of RDMA (**iSCSI/iSER**, NFS, DAFS, SDP, databases, clustering, ... www.clustercomp.org/cluster2006/rdma_and_rait.htm - 15k - Cached - Similar pages

[Remote direct memory access enabled network interface controller ...](#)

A **remote direct memory access** (RDMA) read work request provides a **memory** The last field of the message is a well-known, **iSCSI**, 32 bit Cyclical ... www.freepatentsonline.com/6721806.html - 137k - Cached - Similar pages

[draft-ietf-rddp-rdmap-07 - A Remote Direct Memory Access Protocol ...](#)

Remote Direct Memory Access Write - An RDMA Write operation uses an RDMA a common profile of IPsec to be used with **iSCSI** and the RDDP protocols. ... tools.ietf.org/html/draft-ietf-rddp-rdmap - 179k - Cached - Similar pages

Sponsored Links

[Affordable iSCSI Storage](#)

Full Featured **iSCSI** Appliances Up to 36TB starting at \$2,995 www.celeros.com

[RDMA Enabled Interfaces](#)

GigE TOE and Fibre Channel Low latency and High Throughput www.criticalio.com

[iSCSI -Editor's Choice](#)

STONEFLY -Pioneer of **iSCSI** 2-300TB Snapshot, Replication, Mirroring www.iSCSI.com

remote access direct memory Content at ZDNet UK

Whitepapers This paper gives an overview of **Remote Direct Memory Access** (RDMA) ...
Recently, several protocols such as iSER (iSCSI Extension for RDMA) and ...
www.zdnet.co.uk/tsearch/remote+access+direct+memory.htm - 44k -
[Cached](#) - [Similar pages](#)

direct memory access rdma Content at ZDNet UK

News These are the basic ideas behind **Remote Direct Memory Access**, RDMA, ...
Recently, several protocols such as iSER (iSCSI Extension for RDMA) and SCSI ...
www.zdnet.co.uk/tsearch/direct+memory+access+rdma.htm - 38k -
[Cached](#) - [Similar pages](#)

[1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Download [Google Pack](#): free essential software for your PC

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)